Once we have $F(\mu)$, we can recover $f(t)$ by using the inverse Fourier transform:

$$f(t) = \int_{-\infty}^{\infty} F(\mu) e^{j2\pi\mu t} d\mu . \quad (4.3\text{-}9)$$

Equations (4.3-7) through (4.3-9) have proved theoretically that it is possible to recover a band-limited function from samples of the function obtained at a rate exceeding twice the highest frequency content of the function. Function $H(\mu)$ is called a lowpass filter.

Aliasing

What happens if a band-limited function is sampled at a rate that is less than twice its highest frequency? Figure 4.9 shows the under-sampled case.

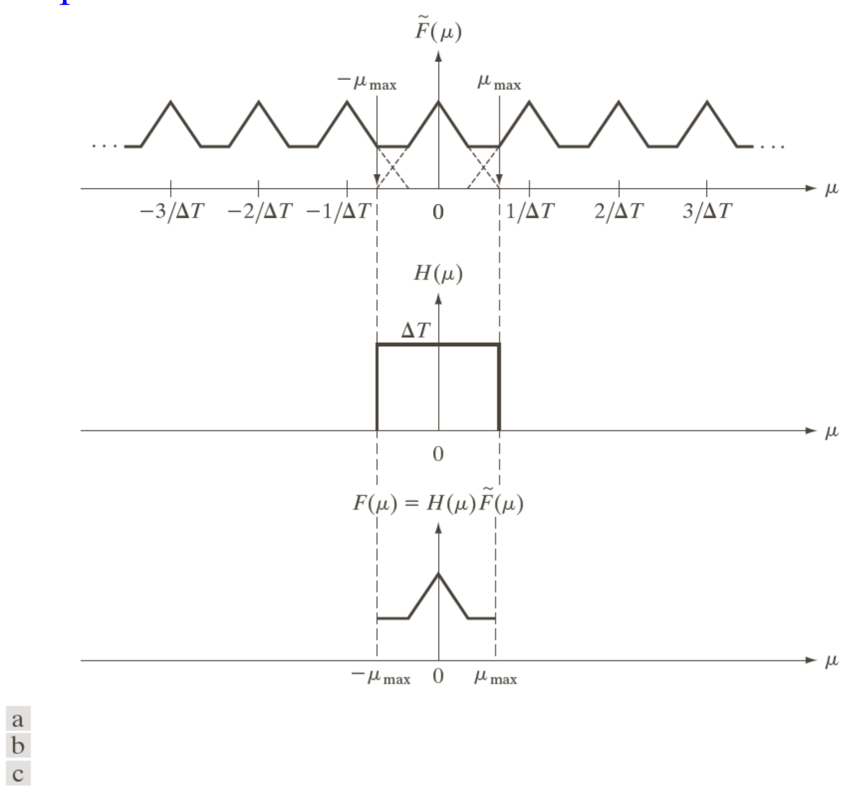

FIGURE 4.9 (a) Fourier transform of an under-sampled, band-limited function. (Interference from adjacent periods is shown dashed in this figure). (b) The same ideal lowpass filter used in Fig. 4.8(b). (c) The product of (a) and (b). The interference from adjacent periods results in aliasing that prevents perfect recovery of $F(\mu)$ and, therefore, of the original, band-limited continuous function. Compare with Fig. 4.8.

The inverse transform would then yield a corrupted function of t. This effect, caused by under-sampling a function, is known as frequency aliasing or simply as aliasing.

Unfortunately, except for some special cases, aliasing is always present in sampled signals. Even if the original sampled function is band-limited, infinite frequency components are introduced the moment we limit the duration of the function.

Suppose we want to limit the duration of a band-limited function $f(t)$ to an interval $[0, T]$. We can do this by multiplying $f(t)$ by the function

$$h(t) = \begin{cases} 1 & 0 \le t \le T \\ 0 & \text{otherwise} \end{cases}, \qquad (4.3\text{-}10)$$

which has the same basic shape as Figure 4.4 (a).

Recall Figure 4.4

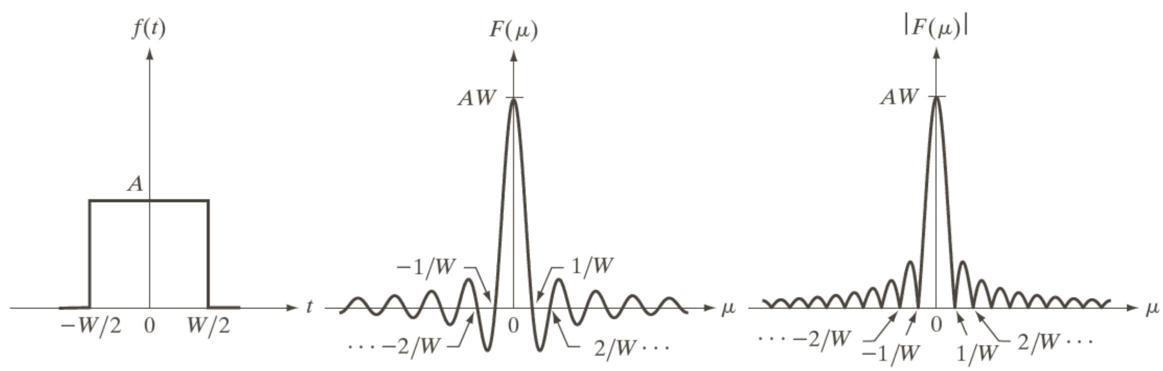

a b c

FIGURE 4.4 (a) A simple function; (b) its Fourier transform; and (c) the spectrum. All functions extend to infinity in both directions.

Even if the transform of $f(t)$ is band-limited, convolving it with $H(\mu)$, will yield a result with frequency components extending to infinity.

Therefore, no function of finite duration can be band-limited. Conversely, a function that is band-limited must extended from $-\infty$ to ∞ .

Although aliasing is an inevitable fact of working with sampled records of finite length, in practice, the effects of aliasing can be reduced by smoothing the input function to attenuate its higher frequencies. This process is called anti-aliasing.

The anti-aliasing must be done before the function is sampled.

Example 4.3: Aliasing

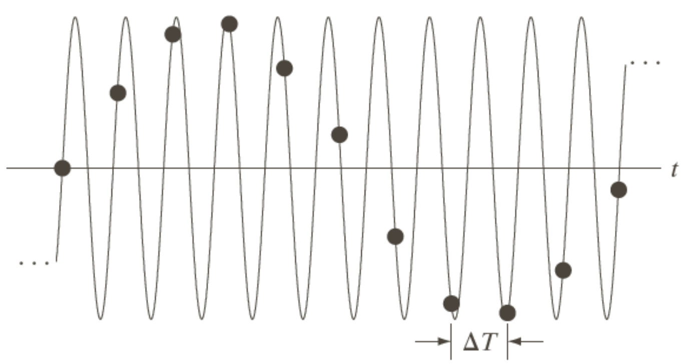

FIGURE 4.10 Illustration of aliasing. The under-sampled function (black dots) looks like a sine wave having a frequency much lower than the frequency of the continuous signal. The period of the sine wave is 2 s, so the zero crossings of the horizontal axis occur every second. ΔT is the separation between samples.

Figure 4.10 shows a classic illustration of aliasing.

Function Reconstruction (Recovery) from Sampled Data

Using the convolution theorem, we can obtain the equivalent result in the spatial domain. From

$$F(\mu) = H(\mu)\tilde{F}(\mu)\,, \tag{4.3-8}$$

it follows that

$$\begin{aligned} f(t) &= \mathscr{F}^{-1}\{F(\mu)\} \\ &= \mathscr{F}^{-1}\left\{H(\mu)\tilde{F}(\mu)\right\} \\ &= h(t)\star\tilde{f}(t)\,. \end{aligned} \tag{4.3-11}$$

It can be shown that substituting (4.3-1) for $\tilde{f}(t)$ into (4.3-11) and using (4.2-20) leads to the following spatial domain expression:

$$f(t) = \sum_{n=-\infty}^{\infty} f(n\Delta T)\mathrm{sinc}\left[(t - n\Delta T)/\Delta T\right]. \tag{4.3-12}$$

Equation (4.3-12) requires an infinite number of terms for the interpolations between samples.

4.4 Extension to Functions of Two Variables

The 2-D Impulse and Its Sifting Property

The impulse, $\delta(t,z)$, of two continuous variables, t and z, is defined as

$$\delta(t,z) = \begin{cases} \infty & \text{if } t = z = 0 \\ 0 & \text{otherwise} \end{cases} \qquad (4.5\text{-}1a)$$

and

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \delta(t,z)dtdz = 1. \qquad (4.5\text{-}1b)$$

The 2-D impulse exhibits the sifting property under integration

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(t,z)\delta(t,z)dtdz = f(0,0), \qquad (4.5\text{-}2)$$

or, more generally for an impulse located at (t_0, z_0),

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(t,z)\delta(t-t_0,z-z_0)dtdz = f(t_0,z_0). \qquad (4.5\text{-}3)$$

For discrete variables x and y, the 2-D discrete impulse is defined as

$$\delta(x,y) = \begin{cases} 1 & \text{if } x = y = 0 \\ 0 & \text{otherwise} \end{cases} \qquad (4.5\text{-}4)$$

and its sifting property is

$$\sum_{x=-\infty}^{\infty}\sum_{y=-\infty}^{\infty} f(x,y)\delta(x,y) = f(0,0). \qquad (4.5\text{-}5)$$

For an impulse located at (x_0, y_0), as shown in Figure 4.12, the sifting property is

$$\sum_{x=-\infty}^{\infty} \sum_{y=-\infty}^{\infty} f(x,y)\delta(x-x_0, y-y_0) = f(x_0, y_0). \qquad (4.5\text{-}6)$$

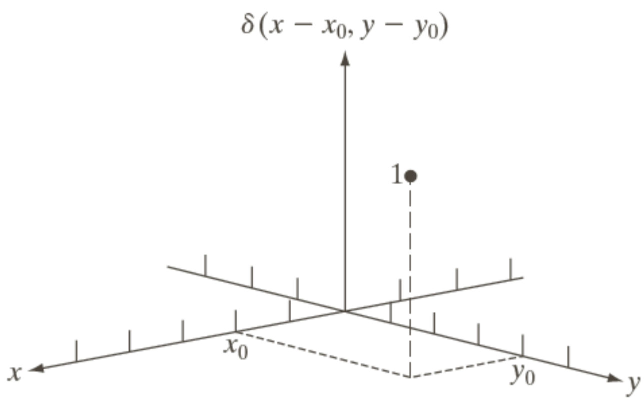

FIGURE 4.12 Two-dimensional unit discrete impulse. Variables x and y are discrete, and δ is zero everywhere except at coordinates (x_0, y_0).

The 2-D Continuous Fourier Transform Pair

Let $f(t,z)$ be a continuous function of two continuous variables. The 2-D continuous Fourier transform pair is given by the expressions

$$F(\mu,\nu) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(t,z)e^{-j2\pi(\mu t+\nu z)}dtdz \qquad (4.5\text{-}7)$$

and

$$f(t,z) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(\mu,\nu)e^{j2\pi(\mu t+\nu z)}d\mu d\nu \qquad (4.5\text{-}8)$$

where μ and ν are the frequency variables.

Example 4.5: Obtaining the 2-D Fourier transform of a simple function

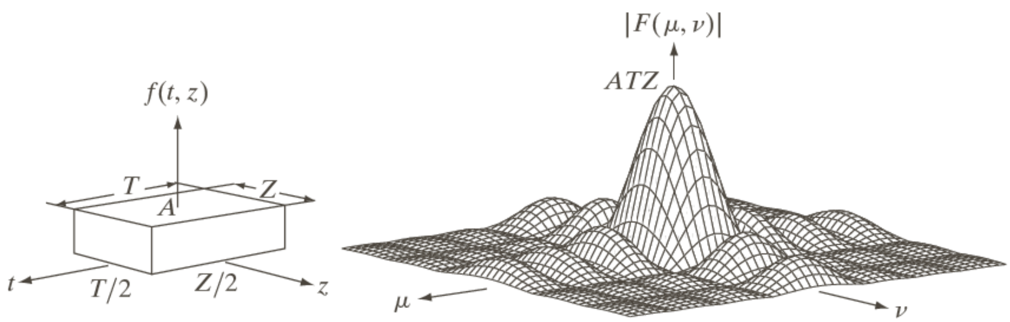

a b

FIGURE 4.13 (a) A 2-D function, and (b) a section of its spectrum (not to scale). The block is longer along the t-axis, so the spectrum is more "contracted" along the μ-axis. Compare with Fig. 4.4.

$$\begin{aligned} F(\mu,\nu) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(t,z)e^{-j2\pi(\mu t+\nu z)}dtdz \\ &= \int_{-T/2}^{T/2}\int_{-Z/2}^{Z/2} Ae^{-j2\pi(\mu t+\nu z)}dtdz \\ &= ATZ\left[\frac{\sin(\pi\mu T)}{(\pi\mu T)}\right]\left[\frac{\sin(\pi\nu Z)}{(\pi\nu Z)}\right] \end{aligned}$$

The magnitude is given by the expression

$$|F(\mu,\nu)| = ATZ\left|\frac{\sin(\pi\mu T)}{(\pi\mu T)}\right|\left|\frac{\sin(\pi\nu Z)}{(\pi\nu Z)}\right|$$

Two-Dimensional Sampling and Sampling Theorem

Sampling in 2-D can be modeled using the sampling function (2-D impulse train):

$$s_{\Delta T \Delta Z}(t, z) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} \delta(t - m\Delta T, z - n\Delta Z) \tag{4.5-9}$$

where ΔT and ΔZ are the separations between samples along the *t*-axis and *z*-axis. Equation (4.5-9) describes a set of the periodic impulses extending along the two axes, as shown in Figure 4.14.

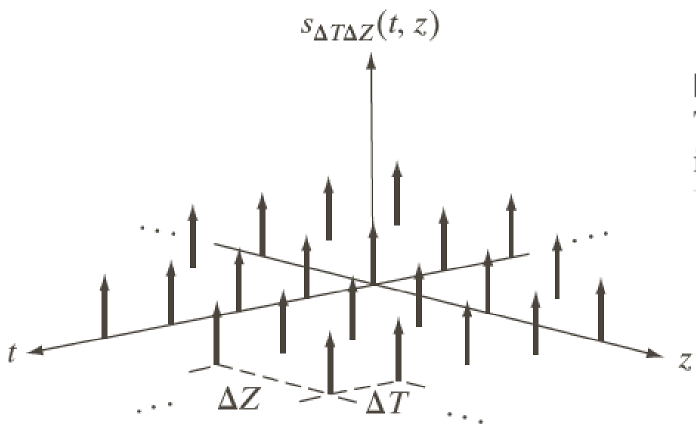

FIGURE 4.14 Two-dimensional impulse train.

Function $f(t, z)$ is said to be band-limited if its Fourier transform is 0 outside a rectangle established by the intervals $[-\mu_{\max}, \mu_{\max}]$ and $[-\nu_{\max}, \nu_{\max}]$:

$$F(\mu, \nu) = 0 \quad \text{for } |\mu| \geq \mu_{\max} \quad \text{and} \quad |\nu| \geq \nu_{\max} \tag{4.5-10}$$

The two-dimensional sampling theorem states that a continuous, band-limited function $f(t, z)$ can be recovered with no error from a set of its samples if the sampling intervals are

$$\Delta T < \frac{1}{2\mu_{\max}} \tag{4.5-11}$$

and

$$\Delta Z < \frac{1}{2\nu_{\max}}, \quad (4.5\text{-}12)$$

or expressed in terms of the sampling rate, if

$$\frac{1}{\Delta T} > 2\mu_{\max} \quad (4.5\text{-}13)$$

and

$$\frac{1}{\Delta Z} > 2\nu_{\max}. \quad (4.5\text{-}14)$$

Figure 4.15 shows the 2-D equivalents of Figure 4.6 (b) and (d).

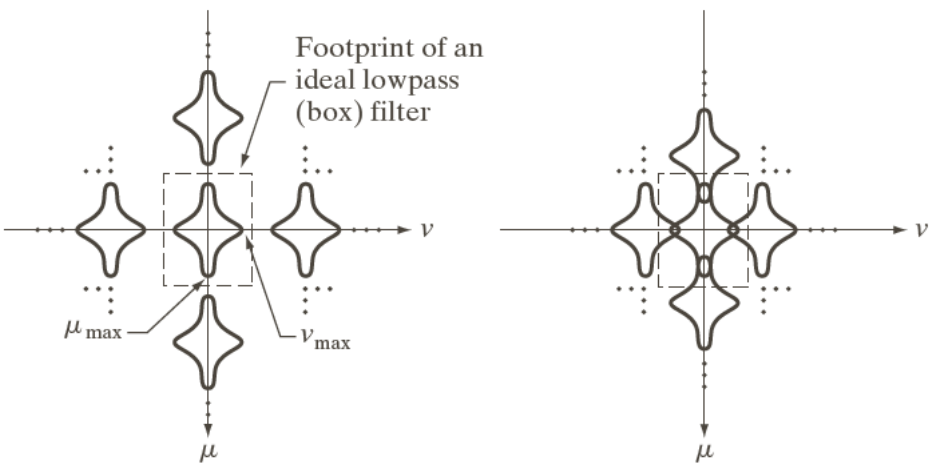

a b

FIGURE 4.15 Two-dimensional Fourier transforms of (a) an over-sampled, and (b) under-sampled band-limited function.

The dashed portion of Figure 4.15 (a) shows the location of an ideal box filter, as shown in Figure 4.13 (a), to achieve the necessary isolation of a single period of the transform for reconstruction of a band-limited function from its samples.

Aliasing in Images

Extension from 1-D aliasing

As in the 1-D case, $f(t, z)$ can be band-limited in general only if it extends infinitely in both coordinates directions. So, aliasing is always present in digital images.

There are two principal manifestations of aliasing in images: spatial aliasing, which is due to under-sampling; and temporal aliasing, which is related to time intervals between images in a sequence of images.

Example 4.6: Aliasing in images

Assume that we have a perfect imaging system that is noiseless and produces an exact digital image of what it sees, but the number of samples it can take is fixed at 96×96 pixels.

If we use this system to digitize checkerboard patterns, it will be able to resolve patterns that are up to 96×96 squares, in which the size of each square is 1×1 pixels.

We are interested in examining what happens when the imaging system is asked to digitize checkerboard patterns that have more than 96×96 squares in the field of view.

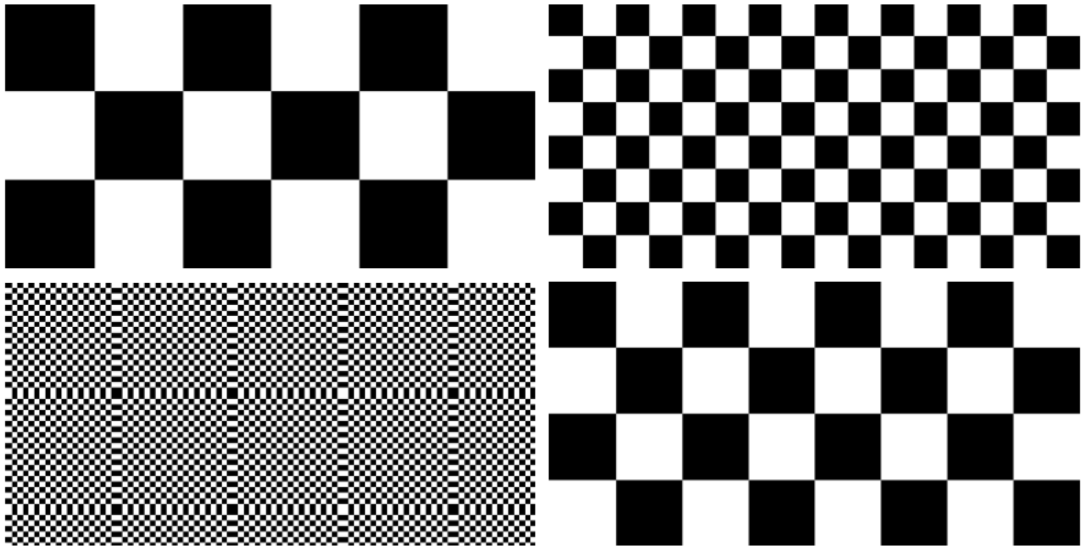

FIGURE 4.16 Aliasing in images. In (a) and (b), the lengths of the sides of the squares are 16 and 6 pixels, respectively, and aliasing is visually negligible. In (c) and (d), the sides of the squares are 0.9174 and 0.4798 pixels, respectively, and the results show significant aliasing. Note that (d) masquerades as a "normal" image.

When the size of the squares is reduced to slightly less than one camera pixel, the system produce a severely aliased image, as Figure 4.16 (c) shows.

When the size of the squares is reduced to slightly less than half of a camera pixel, the system will produce a severely aliased and image, which is very misleading, as Figure 4.16 (d) shows.

The effects of aliasing can be reduced by slightly defocusing the scene to be digitized so that high frequencies are attenuated.

The anti-aliasing filtering has to be done at the "front-end", before the image is sampled. There are no such things as after-the-fact software anti-aliasing filters that can be used to reduce the effects of aliasing caused by violations of the sampling theorem.

Image interpolation and resampling

As in the 1-D case, perfect reconstruction of a band-limited image function from a set of its samples requires 2-D convolution in the spatial domain with a sinc function.

As mentioned previously, this theoretically perfect reconstruction requires interpolation using infinite summations. Therefore, in practice, we need to look for approximations.

One of the most common applications of 2-D interpolation in image processing is in image resizing (zooming and shrinking).

A special case of nearest neighbour interpolation that ties with over-sampling is zooming by pixel replication, which is applicable when we want to increase the size of an image by an integer number of times.

For example, to double the size of an image, we duplicate each column, and then each row of the enlarged image. The intensity-level assignment of each pixel is predetermined by the fact that new locations are exact duplicates of old locations.

Image shrinking is done in a similar manner. Under-sampling is achieved by row-column deletion.

Example 4.7: Illustration of aliasing in resampled images

The effects of aliasing generally are worsened when the size of a digital image is reduced.

a b c

FIGURE 4.17 Illustration of aliasing on resampled images. (a) A digital image with negligible visual aliasing. (b) Result of resizing the image to 50% of its original size by pixel deletion. Aliasing is clearly visible. (c) Result of blurring the image in (a) with a 3×3 averaging filter prior to resizing. The image is slightly more blurred than (b), but aliasing is not longer objectionable. (Original image courtesy of the Signal Compression Laboratory, University of California, Santa Barbara.)

Figure 4.17 (a) is an image purposely created to show the effects of aliasing. Although there are the thinly-spaced parallel lines in all garments worn, there are no objectionable artifacts in Figure 4.17 (a), indicating that the sampling rate was sufficient to avoid visible aliasing.

In Figure 4.17 (b), the image was reduced to 50% of its original size using row-column deletion. The effects of aliasing are quite visible.

Figure 4.17 (c) shows the result of smoothing the image in Figure 4.17 (a) is a 3×3 averaging filter before reducing its size. The improvement over Figure 4.17 (a) is evident.

When we work with images that have strong edge content, the effects of aliasing are seen as block-like image components, called jaggies.

Example 4.8: Illustration of jaggies in image shrinking

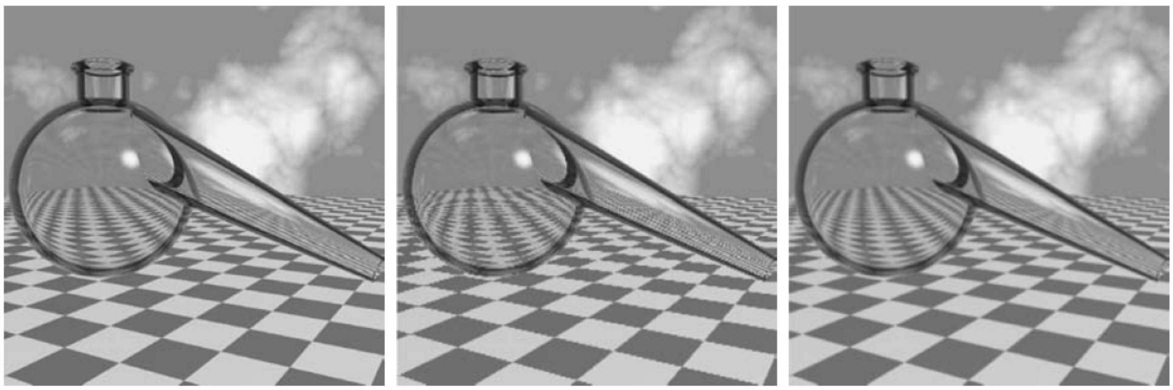

a b c

FIGURE 4.18 Illustration of jaggies. (a) A 1024 × 1024 digital image of a computer-generated scene with negligible visible aliasing. (b) Result of reducing (a) to 25% of its original size using bilinear interpolation. (c) Result of blurring the image in (a) with a 5 × 5 averaging filter prior to resizing it to 25% using bilinear interpolation. (Original image courtesy of D. P. Mitchell, Mental Landscape, LLC.)

Figure 4.18 (a) shows a 1024×1024 image of a computer-generated scene in which aliasing is negligible.

Figure 4.18 (b) is the result of reducing the size by 75% to 256×256 using bilinear interpolation and then using pixel replication to bring the image back to its original size in order to make the effects of aliasing more visible.

Figure 4.18 (c) is the result of using a 5×5 averaging filter prior to reducing the size of image. Compared to Figure 4.18 (b), jaggies in Figure 4.18 (c) were reduced significantly.

Example 4.9: Illustration of jaggies in image zooming

Figure 4.19 (a) shows a 1024×1024 zoomed image generated by pixel replication from a 256×256 section out of the center of Figure 4.18 (a).

Note the "blocky" edges in Figure 4.19 (a).

The zoomed image in Figure 4.19 (b) was generated from the same 256×256 section, but using bilinear interpolation.

The edges in Figure 4.19 (b) are considerably smooth.

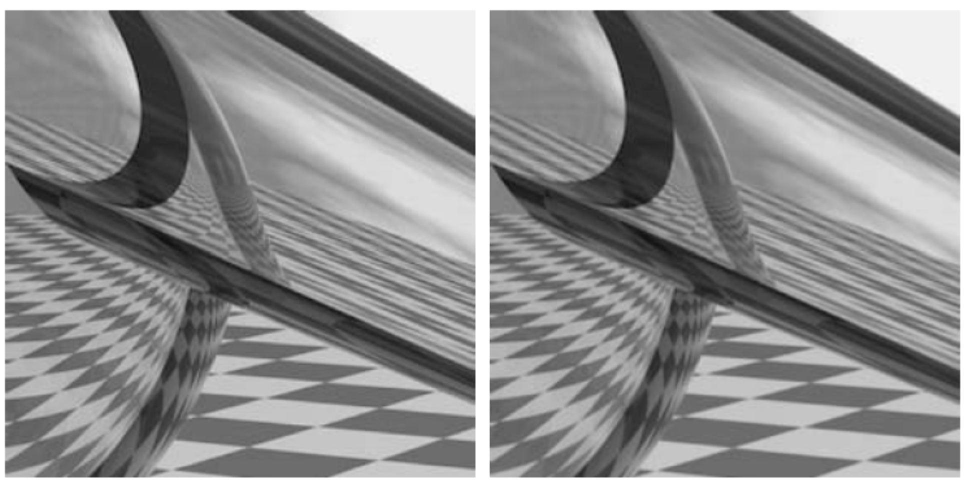

a b

FIGURE 4.19 Image zooming. (a) A 1024×1024 digital image generated by pixel replication from a 256×256 image extracted from the middle of Fig. 4.18(a). (b) Image generated using bi-linear interpolation, showing a significant reduction in jaggies.

4.5 The Discrete Fourier Transform (DFT) of One Variable

The material up to this point may be viewed as the foundation of introducing the discrete Fourier transform (DFT), which is the main part of Chapter 4.

Obtaining the DFT from the Continuous Transform of a Sampled Function

Referring to

$$F(\mu) = \int_{-\infty}^{\infty} f(t)e^{-j2\pi\mu t}dt \qquad (4.2\text{-}16)$$

and by substituting

$$\tilde{f}(t) = \sum_{n=-\infty}^{\infty} f(t)\delta(t - n\Delta T) \qquad (4.3\text{-}1)$$

for $\tilde{f}(t)$, we obtain

$$\begin{aligned} \tilde{F}(\mu) &= \int_{-\infty}^{\infty} \tilde{f}(t)e^{-j2\pi\mu t}dt \qquad (4.4\text{-}1) \\ &= \int_{-\infty}^{\infty} \sum_{n=-\infty}^{\infty} f(t)\delta(t - n\Delta T)e^{-j2\pi\mu t}dt \\ &= \sum_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(t)\delta(t - n\Delta T)e^{-j2\pi\mu t}dt \\ &= \sum_{-\infty}^{\infty} f_n e^{-j2\pi\mu n\Delta T} \qquad (4.4\text{-}2) \end{aligned}$$

Although f_n is a discrete function, its Fourier $\tilde{F}(\mu)$ is continuous and infinitely periodic with period $1/\Delta T$. Therefore, all we need to characterize $\tilde{F}(\mu)$ is one period, and sampling one period is the basis for the DFT.

Suppose that we want to have M equally spaced samples of $\tilde{F}(\mu)$ over the period $\mu = 0$ to $\mu = 1/\Delta T$, we will take the samples at the frequencies

$$\mu = \frac{m}{M\Delta T} \quad m = 0,1,2,\ldots,M-1 \qquad (4.4\text{-}3)$$

Substituting into (4.4-2) and using F_m to denote the result

$$F_m = \sum_{n=0}^{M-1} f_n e^{-j2\pi mn/M} \quad m = 0,1,2,\ldots,M-1 \qquad (4.4\text{-}4)$$

This is the discrete Fourier transform we are seeking.

Conversely, given $\{F_m\}$, we can recover the sample set $\{f_n\}$ by using the inverse discrete Fourier transform (IDFT)

$$f_n = \frac{1}{M}\sum_{m=0}^{M-1} F_m e^{j2\pi mn/M} \quad n = 0,1,2,\ldots,M-1 \qquad (4.4\text{-}5)$$

Equations (4.4-4) and (4.4-5) constitute a discrete Fourier transform pair.

The forward and inverse Fourier transforms exist for any set of samples whose values are finite.

Note that neither expression depends on the sampling interval ΔT nor on the frequency intervals of μ. Therefore, the DFT pair is applicable to any finite set of discrete samples taken uniformly.

For image processing, we intend to use the notation x and y for image coordinates variables and u and v for frequency variables.

Then, (4.4-4) and (4.4-5) become

$$F(u) = \sum_{n=0}^{M-1} f(x)e^{-j2\pi ux/M} \quad u = 0,1,2,\ldots,M-1 \tag{4.4-6}$$

and

$$f(x) = \frac{1}{M}\sum_{u=0}^{M-1} F(u)e^{j2\pi ux/M} \quad x = 0,1,2,\ldots,M-1. \tag{4.4-7}$$

It can be shown that both the forward and inverse discrete transforms are infinitely periodic, with period M,

$$F(u) = F(u + kM) \tag{4.4-8}$$

and

$$f(x) = f(x + kM) \tag{4.4-9}$$

where k is an integer.

The discrete equivalent of convolution in

$$f(t)\star h(t) = \int_{-\infty}^{\infty} f(\tau)h(t-\tau)d\tau \tag{4.2-20}$$

is

$$f(x)\star h(x) = \sum_{m=0}^{M-1} f(m)h(x-m) \tag{4.4-10}$$

for $x = 0,1,2,\ldots,M-1$.

Equation (4.4-10) gives one period of the periodic convolution, therefore, the process is referred to as circular convolution and is a direct result of the periodicity of the DFT and its inverse.

Relationship between the Sampling and Frequency Intervals

If $f(x)$ consists of M samples of a function $f(t)$ taken ΔT units apart, the duration is

$$T = M\Delta T. \qquad (4.4\text{-}11)$$

In the discrete frequency domain, the corresponding spacing, $\Delta\mu$, is

$$\Delta\mu = \frac{1}{M\Delta T} = \frac{1}{T}. \qquad (4.4\text{-}12)$$

The entire frequency range spanned by the M components of the DFT is

$$\Omega = M\Delta\mu = \frac{1}{\Delta T}. \qquad (4.4\text{-}13)$$

The resolution in frequency, $\Delta\mu$, of the DFT depends on the duration T over which the continuous function, $f(t)$, is sampled. The range of frequencies spanned by the DFT depends on the sampling interval ΔT.

Example 4.4: The mechanics of computing the DFT

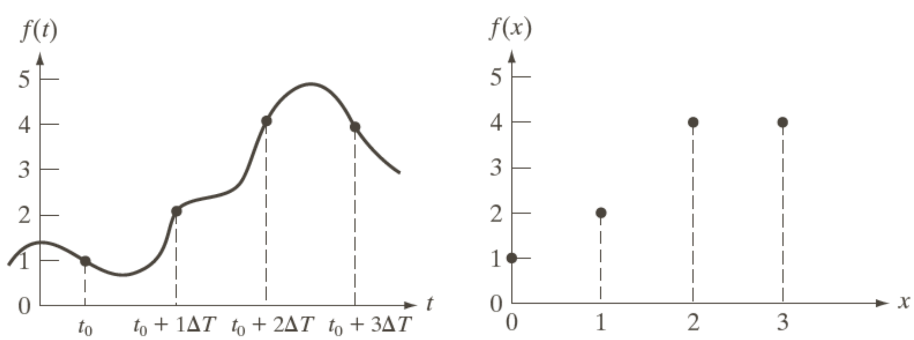

a b
FIGURE 4.11 (a) A function, and (b) samples in the x-domain. In (a), t is a continuous variable; in (b), x represents integer values.

Figure 4.11 (a) shows four samples of a continuous function $f(t)$, and Figure 4.11 (b) shows the sampled values in the x-domain.

From

$$F(u) = \sum_{n=0}^{M-1} f(x) e^{-j2\pi ux/M} \quad u = 0,1,2,\ldots,M-1, \tag{4.4-6}$$

we obtain

$$\begin{aligned} F(0) &= \sum_{x=0}^{3} f(x) = [f(0) + f(1) + f(2) + f(3)] \\ &= 1 + 2 + 4 + 4 = 11, \end{aligned}$$

and the next value of $F(u)$

$$\begin{aligned} F(1) &= \sum_{x=0}^{3} f(x) e^{-j2\pi(1)x/4} \\ &= 1e^{0} + 2e^{-j\pi/2} + 4e^{-j\pi} + 4e^{-j3\pi/2} = -3 + 2j \end{aligned}$$

Similarly, we have $F(2) = -(1 + 0j)$ and $F(3) = -(3 + 2j)$.

If we were given $F(u)$ and were asked to compute its inverse, we would use the inverse transform

$$f(x) = \frac{1}{M} \sum_{u=0}^{M-1} F(u) e^{j2\pi ux/M} \quad x = 0,1,2,\ldots,M-1. \tag{4.4-7}$$

For example,

$$\begin{aligned} f(0) &= \frac{1}{4} \sum_{u=0}^{3} F(u) e^{j2\pi u(0)} = \frac{1}{4} \sum_{u=0}^{3} F(u) \\ &= \frac{1}{4}[11 - 3 + 2j - 1 - 3 - 2j] = 1 \end{aligned}$$

The 2-D Discrete Fourier Transform and Its Inverse

The 2-D discrete Fourier transform (DFT) is

$$F(u,v) = \sum_{x=0}^{M-1}\sum_{y=0}^{N-1} f(x,y) e^{-j2\pi(ux/M+vy/N)}, \qquad (4.5\text{-}15)$$

for $u = 0,1,2,..,M-1$ and $v = 0,1,2,..,N-1$.

Given $F(u,v)$, we can obtain $f(x,y)$ by using the inverse discrete Fourier transform (IDFT):

$$f(x,y) = \frac{1}{MN}\sum_{u=0}^{M-1}\sum_{v=0}^{N-1} F(u,v) e^{j2\pi(ux/M+vy/N)}, \qquad (4.5\text{-}16)$$

for $x = 0,1,2,..,M-1$ and $y = 0,1,2,..,N-1$.

Equations (4.5-15) and (4.5-16) constitute the 2-D discrete Fourier transform pair.